Computing in High Energy Physics

Irwin Gaines DOE/FNAL

HEPAP Meeting 24-Jul-2003

Outline

- Computing as one of many enabling technologies in HEP
- Special features of computing
  - Universality (everyone does it)
  - Connections with real world
- Computing Issues
- Computing Topics

Computing as Enabling Technology

- Computing is one of many tools we use to do our research
- As with other technologies, we are constantly pushing the envelope to get more performance for less (data acquisition systems, emulators, parallel processing farms, commodity processors and storage, ...); often ahead of but sometimes behind (C++) industry

Why software systems aren't just like building a drift chamber

- Everyone uses computers (so we have been reluctant to use specialized tools that dis-enfranchise some users [C++])
- Slow penetration of software engineering discipline (contrast to building a drift chamber – done by experts)
- Much more commonalities with real world (so we can't make our own standards, must compete for manpower, have potential for collaboration outside the field)

Computing Issues

- Getting credit for our computing innovations (we can’t live off the web forever). We push the envelope, but we aren’t sexy.
- Career paths for computing physicists
- Ensuring resources for computing systems (especially for people)
  - Within experiments
  - From funding agencies and OMB
- Interagency, interdisciplinary, and International cooperation/collaboration

Agreement on 5 principles:

- The cost and complexity of 21st Century Science requires the creation of **advanced and coherent global Infostructure** (information infrastructure).
- The construction of a coherent Global Infostructure for Science requires **definition and drivers from Global Applications** (that will also communicate with each other)
- Further, **forefront Information Technology must be incorporated** into this Global Infostructure for the Applications to reach their full potential for changing the way science is done.
- **LHC** is a near term Global Application requiring advanced and un-invented Infostructure and is **ahead in planning** compared to many others.
- U.S. agencies must work together for effective U.S. participation on Global scale infostructure, and the successful execution of the LHC program in a **4 way agency partnership**, ***with international cooperation in view***.

Partnerships

- International: Europe/US/Asia (Europe in particular putting heavy funding into “Grid”)
- Interagency: Different funding agencies
- Interdisciplinary: Application scientists and computer scientists

Moving from era of interoperability (everyone develops their own tools and we figure out how to make them work together) to true collaboration!

International HEPCCC Charter

Preamble

- As the problems posed by computing in HEP projects are becoming more and more global in nature, the HEP community recognizes the need for a global forum for HEP computing. ICFA therefore sponsors I_HEPCCC, a forum to primarily help with an efficient information exchange on computing issues between the major HEP centers in the world.

Missions:

- IHEPCCC primary mission is to act as a forum between the main persons in charge of HEP computing, by gathering and distributing information about all relevant issues in HEP computing, and especially those with a global nature. Typical examples are information exchanges about new technology trends, computing centers strategic policies, security issues, recommendation of standard practices, presentation of R&D results, comparison of various equipments performances.
- The other missions include:
  - Issuing statements and recommendations concerning computing in the HEP community.
  - Serving as an interface to other scientific domains on matters of computing.
  - Working in close connection with the ICFA SCIC, the physics regional organizations, and the HICB coordinating the grid projects in HEP.
  - Reporting to ICFA

Computing Topics

- Simulation as 3rd pillar of scientific discovery (SciDAC program)
- Special purpose lattice gauge computers
- Experimental Computing as a project (LHC experiments)
- Networks
- Grids

Testimony of Dr. Raymond L. Orbach
Director, Office of Science, U.S. Department of Energy
before the U.S. House of Representatives Committee on Science
July 16, 2003 :

“The tools for scientific discovery have changed. Previously, science had been limited to experiment and theory as the two pillars for investigation of the laws of nature. With the advent of what many refer to as "Ultra-Scale" computation," a third pillar-simulation-has been added to the foundation of scientific discovery. Modern computational methods are developing at such a rapid rate that computational simulation is possible on a scale that is comparable in importance with experiment and theory. The remarkable power of these facilities is opening new vistas for science and technology.”

But he did not cite high energy physics.

Scientific Discovery through Advanced Computing (SciDAC)

- Orbach: “To address the need for mathematical and software tools, and to develop highly efficient simulation codes for scientific discovery, the Office of Science launched the Scientific Discovery through Advanced Computing (SciDAC) program. We have assembled interdisciplinary teams and collaborations to develop the necessary state-of-the-art mathematical algorithms and software, supported by appropriate hardware and middleware infrastructure to use terascale computers effectively to advance fundamental scientific research essential to the DOE mission
- HEP SciDAC projects include 2 Supernova simulation projects, accelerator modeling,lattice Gauge calculations, and Particle Physics data Grid Collaboratory.
- Mainstream experimental HEP is not part of these initiatives

National Computational Infrastructure for Lattice Gauge Theory (Sugar- PI)

- Representing ~60 Theorists in the US. Funding to 3 labs and 6 universities
- National effort to regain US competitiveness
  - put in place the software and hardware needed for accurate lattice QCD calculations

National Computational Infrastructure for Lattice Gauge Theory (Sugar- PI)

- Huge strides made in collaborative approach
  - + starting to work with computer scientists on performance metrics and optimization of code
- Accurate computations of important scientific constants requires tens of Tflop years
  - Need highly cost-effective Topical Computing Centers for Lattice QCD – aiming at below $1/Mflop and targeting two different machine architectures – (1) Custom built for QCD and (2) Commodity PC Clusters with low latency networking

National Computational Infrastructure for Lattice Gauge Theory (Sugar- PI)

- QCDOC design is complete, custom chips have been delivered, boards being tested now
- Propose to build a 128 node development machine this year, and a full 5-10 Tflop machine next year (UK-QCD and Riken each have already paid for a large machine
- Wilczek review panel recommended to proceed in Feb 2003
- Funds not available! ($2M this year, $5-10M next year)

US ATLAS and CMS Software and Computing Projects

- Both experiments have defined Software & Computing Projects as part of the US LHC Research Program (follow on to LHC construction)
- ~$40M projects each, roughly 2/3 personnel, rest hardware at Tier 1 and Tier 2 regional computing centers
- Detailed resource loaded schedules, milestones, etc
- This has enabled early hiring of significant number of software engineers.
- Funding goes to project managers, not to individual institutions (as construction project is funded)
- Other experiments might benefit from similar arrangements

Centres taking part in the LCG-1

around the world around the clock

LHC Computing Model

- Distributed model from the start (distributed resources + coherent global access to data)
- Must support
  - Production (reconstruction, simulation)
    - Scheduled, predictable, batch
    - Run by experiment or physics group
    - Highly compute intensive, accesses predictable data sets
  - Data Analysis (including calibration and monitoring)
    - Random, chaotic, often interactive
    - Run by individuals and small groups
    - Mostly data intensive, accesses random data
    - Highly collaborative
  - Code development and testing
    - Highly interactive
    - Highly collaborative

LHC Computing Facilities Model

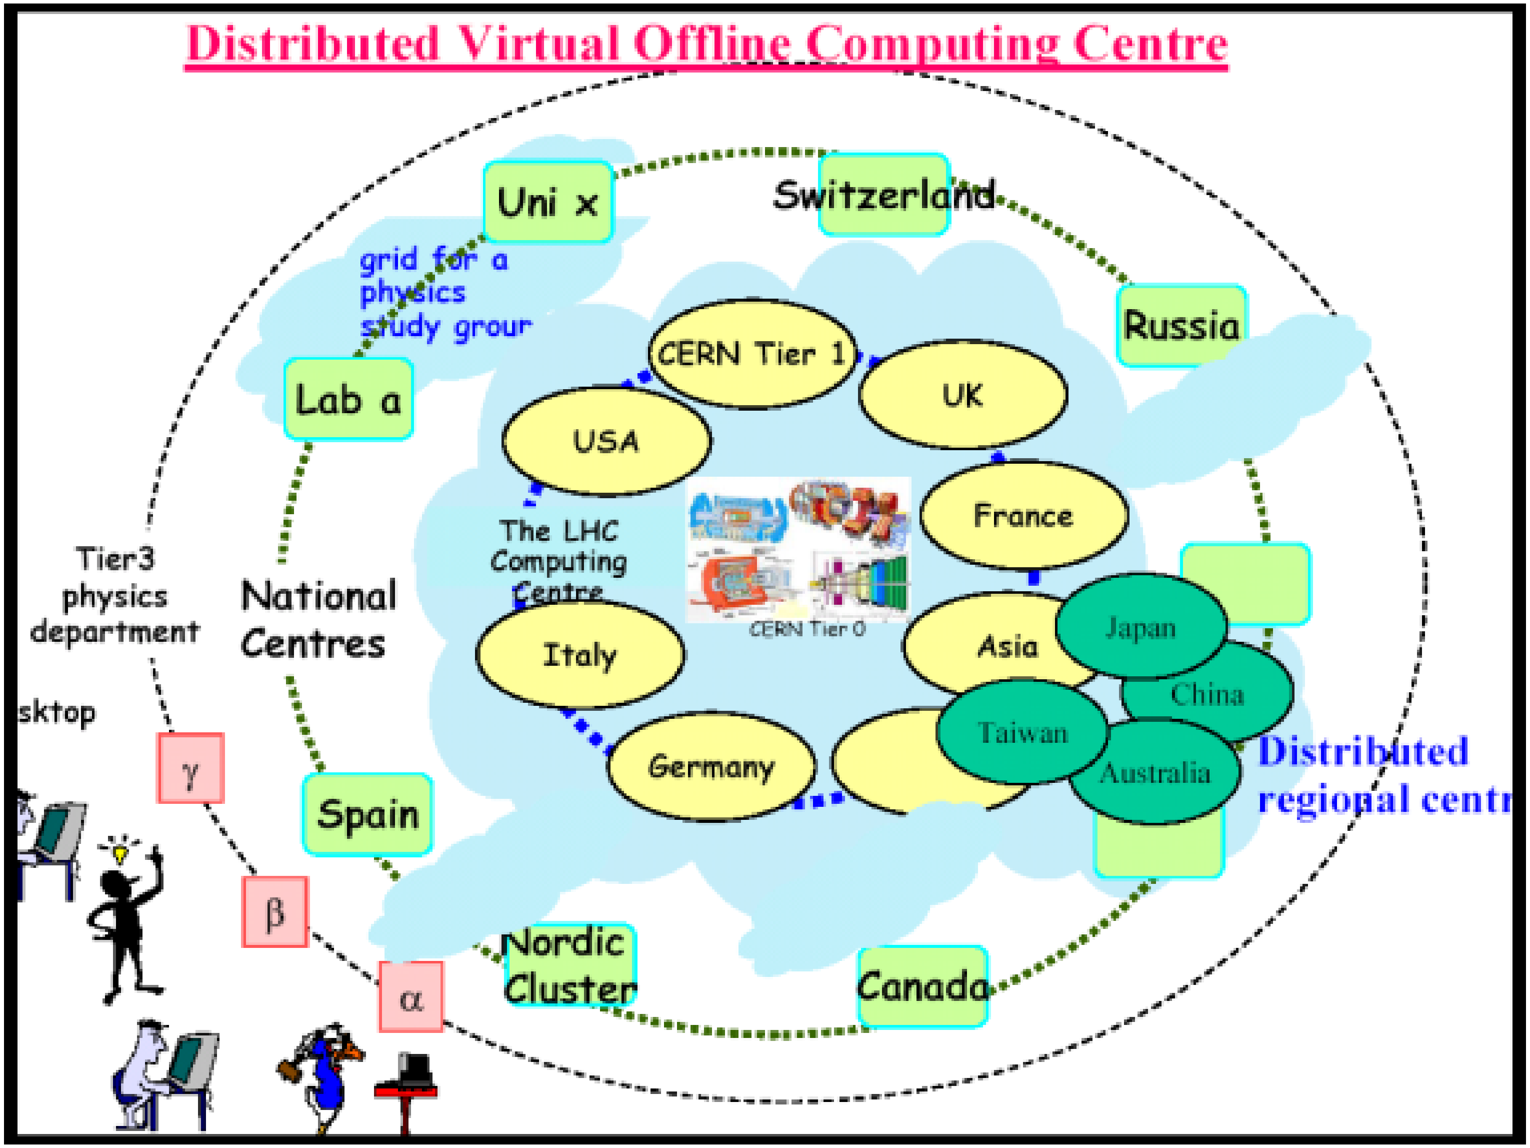

Networking

- We are heavy consumers of network bandwidth
- This will increase dramatically as new generations of experiments accumulate massive amounts of data and develop techniques for distributed data analysis
- Current network usage is strange combination of different networks from multiple funding sources (ESNet for labs, Internet2 for universities, ad-hoc international networks)
- For a long time we enjoyed unsurpassed connectivity, but with deregulation in Europe leading to much lower prices, the European research networks (GEANT) are now at least the equal of ours

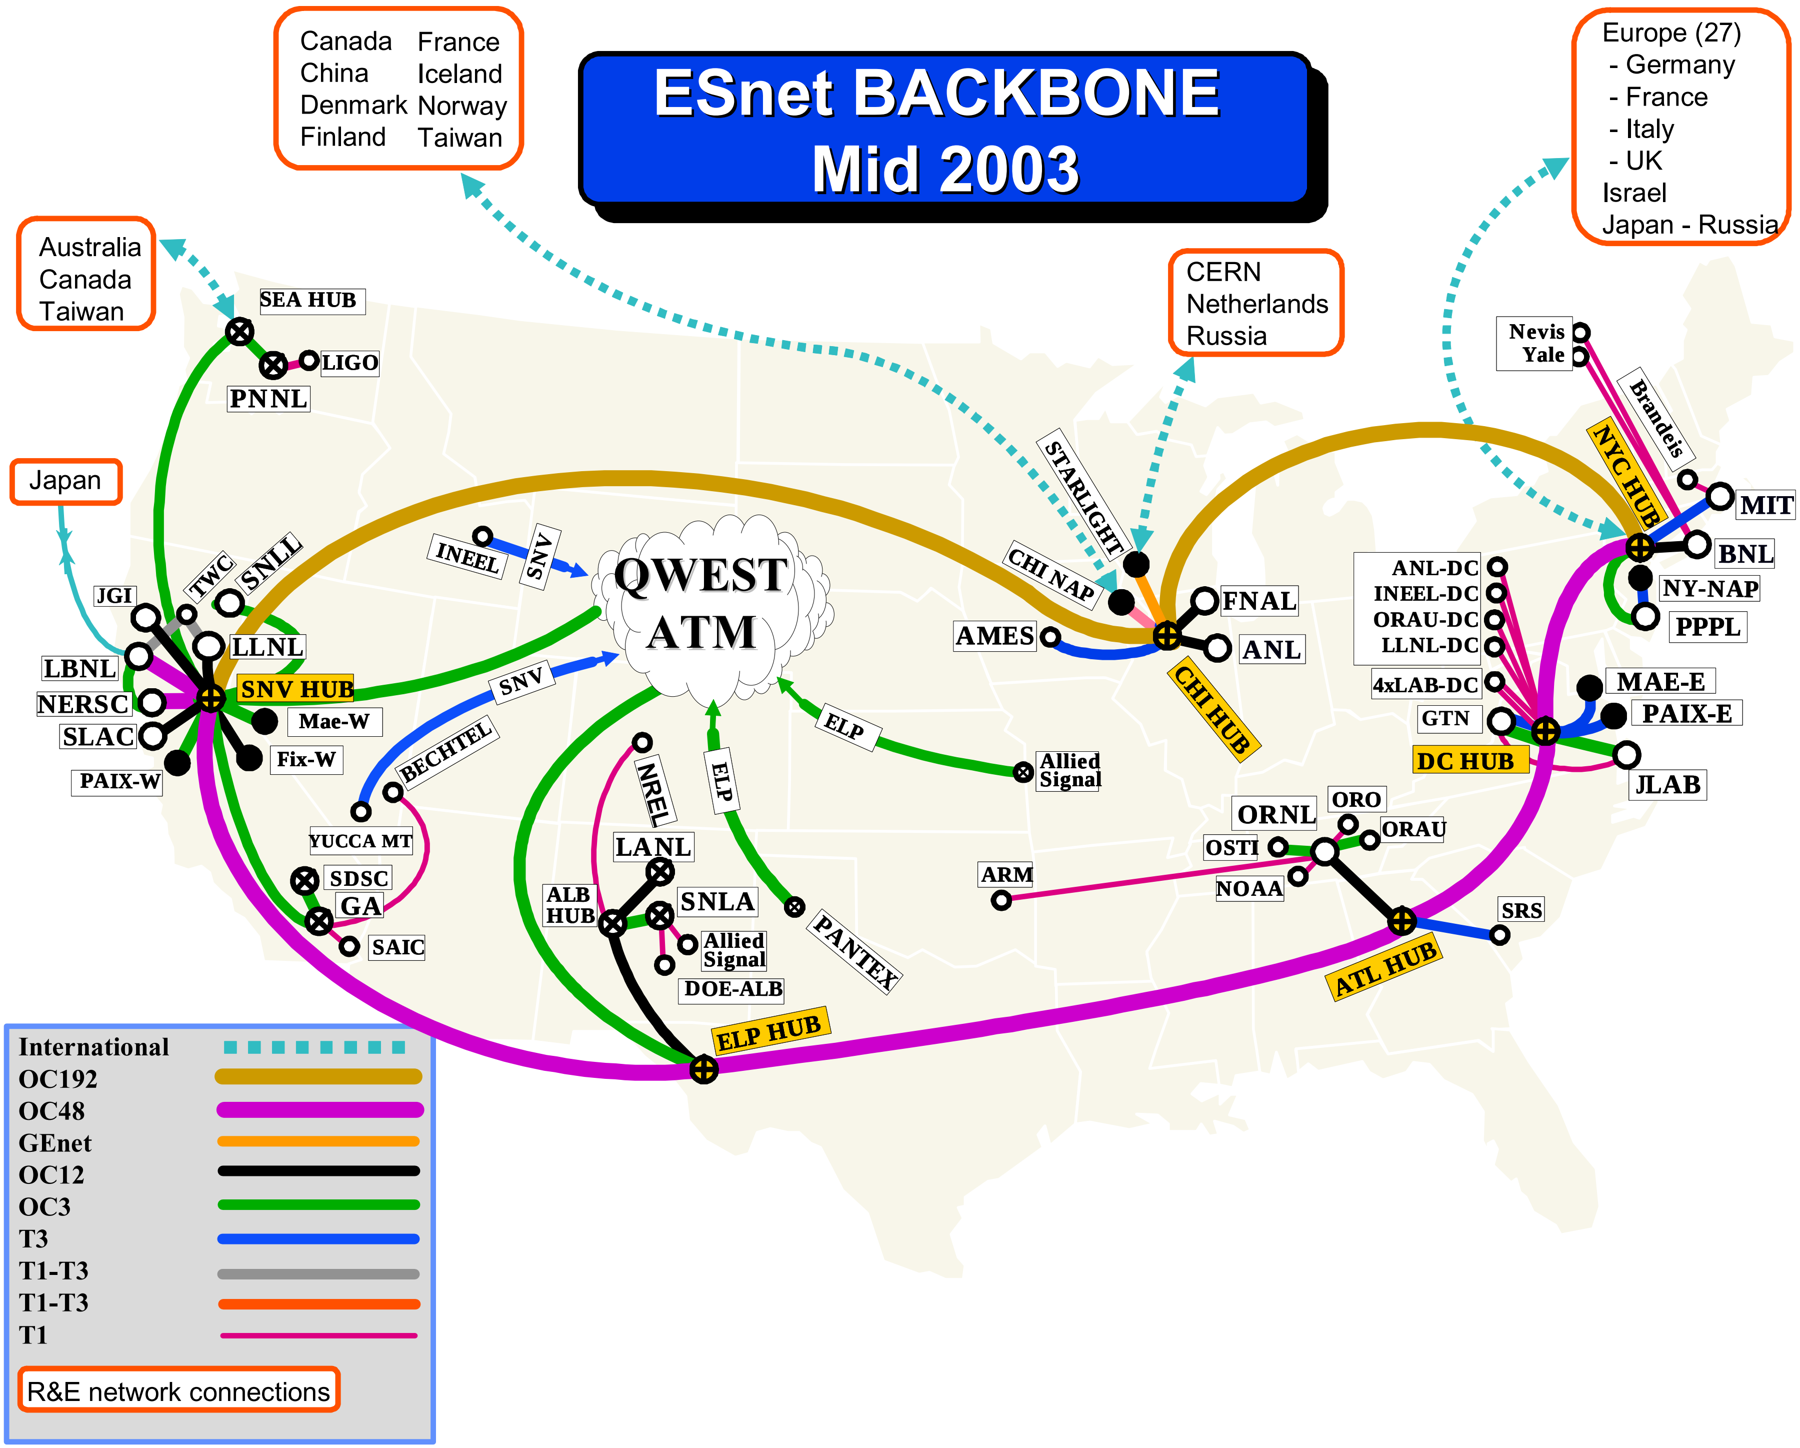
ESnet BACKBONE
Mid 2003
Canada France
China Iceland
Denmark Norway
Finland Taiwan
Australia
Canada
Taiwan
Japan
Europe (27)
- Germany
- France
- Italy
- UK
Israel
Japan - Russia
CERN
Netherlands
Russia
SEA HUB
LIGO
PNNL
JGI
TWC
SNLL
LLNL
LBNL
NERSC
SLAC
PAIX-W
SNV HUB
Mae-W
Fix-W
INEEL
SNV
QWEST ATM
SNV
BECHTEL
YUCCA MT
SDSC
GA
SAIC
NREL
ELP
LANL
ALB HUB
SNLA
Allied Signal
DOE-ALB
PANTEX
ELP HUB
ELP
Allied Signal
STARLIGHT
CHI NAP
AMES
FNAL
ANL
CHI HUB
ARM
ORNL
ORO
OSTI
ORAU
NOAA
SRS
ATL HUB
Nevis
Yale
Brandeis
NYC HUB
MIT
BNL
NY-NAP
PPPL
ANL-DC
INEEL-DC
ORAU-DC
LLNL-DC
4xLAB-DC
GTN
MAE-E
PAIX-E
DC HUB
JLAB
International
OC192
OC48
GEnet
OC12
OC3
T3
T1-T3
T1-T3
T1
R&E network connections

22 May 2003

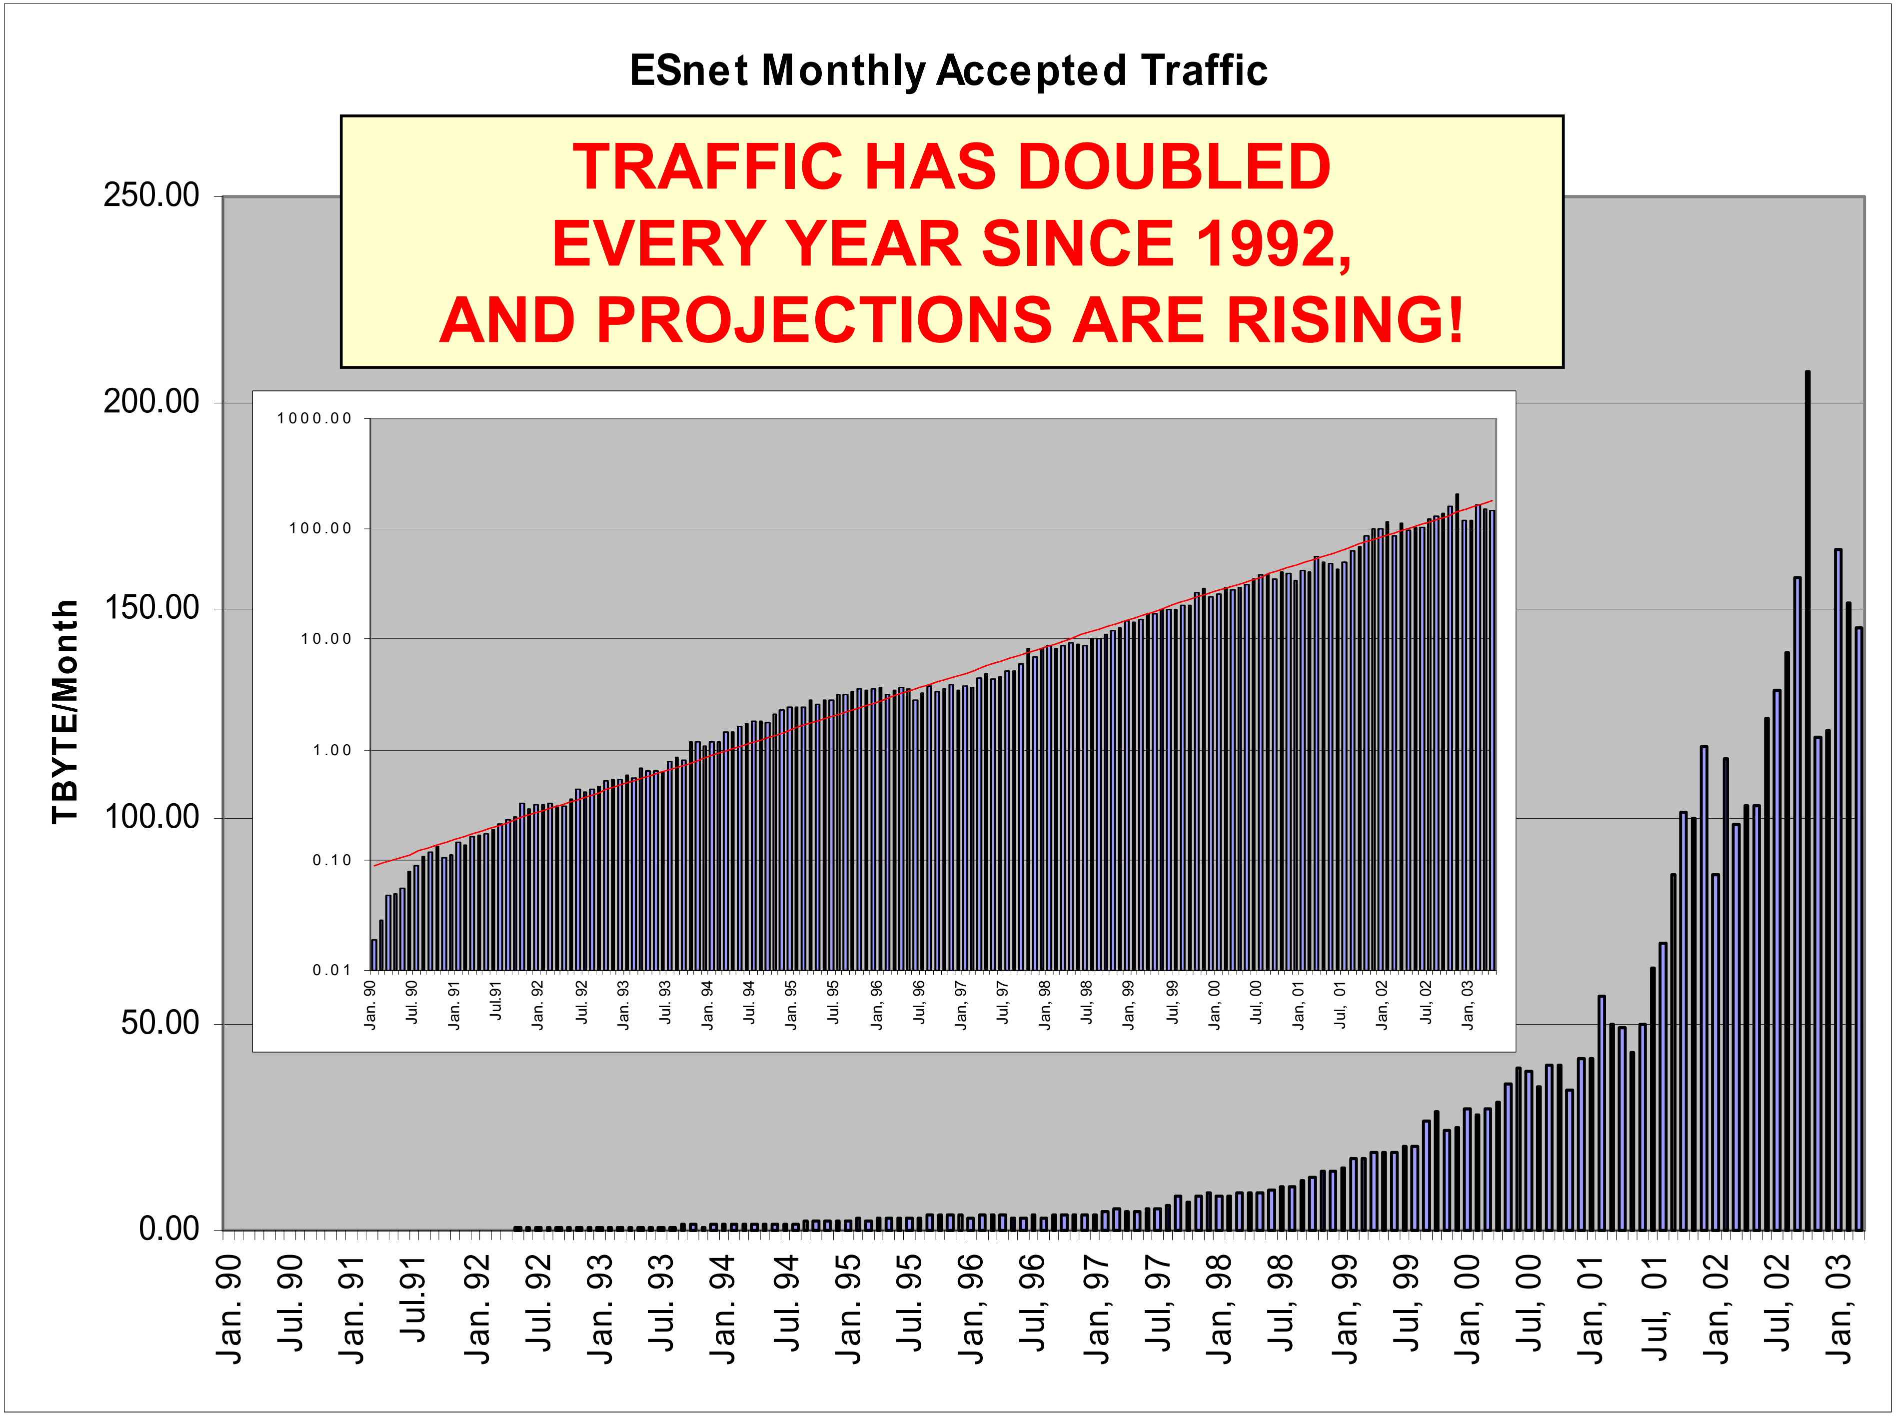
ESnet Monthly Accepted Traffic
TRAFFIC HAS DOUBLED
EVERY YEAR SINCE 1992,
AND PROJECTIONS ARE RISING!
TBYTE/Month
250.00
200.00
150.00
100.00
50.00
0.00
1000.00
100.00
10.00
1.00
0.10
0.01
Jan. 90
Jul. 90
Jan. 91
Jul.91
Jan. 92
Jul. 92
Jan. 93
Jul. 93
Jan. 94
Jul. 94
Jan. 95
Jul. 95
Jan, 96
Jul, 96
Jan, 97
Jul, 97
Jan, 98
Jul, 98
Jan, 99
Jul, 99
Jan, 00
Jul, 00
Jan, 01
Jul, 01
Jan, 02
Jul, 02
Jan, 03

DOE Science Network Workshop

- Held in Reston at beginning of June
- Outlined the case for a three tier network model
  - Production networks
  - High impact networks
  - Research networks
- Substantially higher bandwidth and better end user connectivity than today
- Report available soon, as input to FY05 budget process

Grids

- The grid is much more than just a way to manage a set of distributed computing resources; allows flexible and dynamic use of resources not under your control
- HEP has long made use of distributed computing, and particularly stresses the grid for data intensive applications
- Fully functioning grid will enable analysis paradigms and data access not previously possible
- HEP Grid research projects (GriPhyN, iVDGL, PPDG) have made important contributions to development and deployment of grid software
- Issue is how the grid software will be supported long term and how the production grids will be managed and operated: proposed Grid3 as next step in creating a permanent grid infrastructure

U.S. Grids for Physics: Grid2003 and Beyond

A U.S. Grid Infrastructure for Science

The U.S. Trillium Physics Grid projects have demonstrated the capabilities and benefits of distributed data grids for physics in the U.S. A common set of core Grid technologies based on Globus and Condor and packaged as the Virtual Data Toolkit are now used in the U.S. and European physics grid communities.

The U.S. High Energy and Nuclear Physics community will continue to work closely with grid technology providers to develop and construct the permanent Grid infrastructure necessary to meet the data processing and analysis requirements of the experiments at the Large Hadron Collider. Over the next few years we will build working grids with improved throughput and capabilities. We will continue to collaborate with our European colleagues such as the LHC Computing Grid and the proposed EGEE.

The grid infrastructure for U.S particle physics will provide the backbone for sharing computational and data storage resources. National Laboratories and Universities will provide a common robust infrastructure which supports the data analysis required by the global experiment communities.

Grid3: Purpose and Goals

The Grid3 project is the next step in creating a permanent grid infrastructure for the U.S. physics community. We plan a first demonstration at the SC2003 conference in November. Grid3 is a collaboration between the Trillium Physics Grid projects and the US ATLAS and US CMS software and computing projects. Grid3 will

- Demonstrate a U.S. grid of ~15 sites, used by ~4 communities, with ~2 Terabytes transferred per day and execution of up to 1000 physics simulations a day;
- Integrate and coordinate U.S. grid development efforts to operate with the first deployments of the LHC Computing Grid;
- Provide the next phase of iVDGL as well as a laboratory for computer science applications.

Driving Application: LHC Data Challenges

Each of the LHC experiments is planning a series of data processing challenges of increasing scale and complexity. The challenges demonstrate the capability of the U.S. Grid infrastructure as analysis, workflow, data rate, total throughput, load, stability and the interaction of multiple grid sites. The goal is to ensure that the computing infrastructure and services and a complete suite of software at the participating institutions are ready when the beam turns on.

The Open Science Grid

The Open Science Grid (OSG) is proposed as a follow-up project to Grid3. Under OSG, the U.S. Grid will expand to match the needs of the broader U.S. scientific community. It will collaborate with international grid projects such as the European and Asian science grids. Together, these grids will provide the necessary worldwide computing environment for LHC physics.

Expanding the Community

The U.S. Grid, from Grid3 forward, is expected to become available to scientific communities outside of LHC and particle physics. The Sloan Digital Sky Survey and the Laser Interferometer Gravitational Wave Observatory are already participating in Grid3. We anticipate that academic, commercial, research, and governmental institutions will all be able to exploit the compute, data storage, network, and support resources the grid will provide.

The U.S. Grid will also provide unparalleled educational opportunities for students at all levels, as resources will be made available for outreach activities student-led research projects, including projects at small or remote institutions.

Conclusions

- Make sure we take part in new high performance computing initiatives (even if we are not interested in supercomputers)
- Clearly state the case to establish our computing and networking needs
- Include computing costs in budget planning
- Continue to play leading role in grid development and deployment